

Maps for Verbs¹

Paul Cohen
Department of Computer Science
Lederle GRC
University of Massachusetts at Amherst
Amherst, MA, 01003 USA
phone: 413 545 3638
email: cohen@cs.umass.edu

Abstract: This paper describes a representation of the meanings of verbs based on the dynamics of interactions between two agents or objects. The representation treats interactions as having three phases, before, during and after contact. Maps for these phases are constructed. Trajectories through these maps correspond to different types of interactions and are denoted by different verbs. We summarize the results of experiments on learning and reasoning with maps.

1. Introduction

Much of what we know and say refers to the dynamics of our world. Here I include our mental world, the world of social interactions, and other not-entirely-physical environments. We have a large class of linguistic objects – verbs – devoted entirely to expressing dynamics. Subtle differences in the meanings of verbs, which linguists call “manner,” are also often dynamical. For instance, the difference between “nudge” and “shove” is partly a matter of mass, movement, and energy transfer from one body to another; and partly a matter of intention. Some AI researchers – those concerned with stochastic control, Markov decision processes, qualitative physics and the like – have developed representations of dynamics that machines can reason with. However, the knowledge representation community and ontology engineers seem satisfied with declarative statements *about* dynamics rather than representations *of* dynamics. They say, “Two agents collided and one fell down,” but they don’t describe the collision or the dynamics of falling. Ontologies generally describe everything about movement but the movement itself. Like a dictionary, they tell us that strolling is a casual, unhurried kind of walking, but they don’t represent the actual movement.

Why should ontologies represent dynamics? Dynamical representations are compact in the sense that a single representation can describe dozens of related concepts. They make explicit the manner of movement and thus make fine distinctions between word meanings. They are grounded in the sense that one can attach sensors to a corpus of dynamical concepts and have the corpus recognize concepts from sensed movement –

¹ Acknowledgment: This work benefitted from discussions with Marc Atkin, David Fisher, Greg Jorstad, Wendy Lehnert, Tim Oates, Michael Rosenstein, and David Westbrook. The project is supported by DARPA award number F499620-97-1-0485 to Cohen.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 1998		2. REPORT TYPE		3. DATES COVERED 00-00-1998 to 00-00-1998	
4. TITLE AND SUBTITLE Maps for Verbs				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Massachusetts, Department of Computer Science, Amherst, MA, 01002				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 13	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

something no ontology can currently do (Rosenstein, Cohen, Schmill and Atkin, 1997). Dynamical representations of physical interactions are easily learned from observations of dynamics (Rosenstein et al., 1997) this is true also of dynamical representations of linguistic constructs (e.g., Regier, 1995; Elman, 1995). The strongest reason to consider dynamics as a foundation for ontologies, I think, is that the knowledge of the youngest humans – neonates and infants – is produced by interacting physically with the world. Neonates are capable of movement, but nobody credits them with conceptual thought. Concepts must therefore result from neonatal and infant experience, which is primarily sensorimotor experience. Much of my research is devoted to showing how a sensorimotor agent (a robot) can acquire a conceptual system (i.e., an ontology) through physical interaction with its environment. Dynamical representations are central to this work.

In this paper I sketch a dynamical representation of verb meanings. Parts of the representation have been implemented, as have modes of reasoning with the representation. Research on learning such representations from interaction with the environment is in progress.

One may be tempted to say, "Dynamical representations of verbs makes sense, because verbs denote activities, but surely you aren't suggesting dynamical representations of objects." Not exactly, but I am suggesting that *classes* of objects are differentiated by how we interact with them, that *concepts* are abstractions over those classes, and that *meanings* of concepts are in large part predictive models of how interactions with objects will unfold. Let me illustrate with the photographs in Figure 1. In the *interactionist* view, which is attributed to Lakoff and Johnson (Lakoff, 1984; Lakoff and Johnson, 1980) and to which I subscribe, category distinctions are based on activity. For months, plastic frogs and spoons were functionally indistinguishable to Allegra: She would grasp either, put it in her mouth, and chew. The fact that we consider the frog a *toy*, and a spoon a *utensil* doesn't matter to her. These are adult categories, not infant categories. On the interactionist account, only when Allegra uses the spoon to eat food will she differentiate it from the frog, and only then will she form a category that resembles in its membership those items we adults call "utensils."




Figure 1. Allegra grasps and mouths a frog. Months later she uses a spoon to feed herself.

So much for categories, but what about concepts and meaning? Here I want to point out that except for formal, mathematical objects, many things – perhaps most – are defined in

terms of what we do with them, or how they were formed, or how they behave. One *could* define spoons in volumetric terms, or in terms of the materials from which they are fabricated, but that's not how we think of spoons unless what we're trying to do is design or fabricate spoons, so even in this case the definition is tied to activity. So the concept of spoon is really a representation of the activities spoons are involved in, and the meaning of this concept is essentially *predictive*: What it means to be a spoon is just what happens to spoons in various activities.


One might try again to limit the scope of this interactionist argument – to say, "Even if you can ground physical concepts in dynamics – and it's true that many verbs denote physical action – the meanings of some verbs, such as read, think, give, plan, and so on, have to do with mental, not physical, activities, primarily. Similarly, words like wealth, information, credibility, and so on, denote nonphysical attributes or things. Surely you aren't suggesting a dynamical representation for these concepts, too." Not exactly, but I am taken with Lakoff and Johnson's (1980) argument that metaphor extends physical concepts to nonphysical ones. Indeed, reading, thinking, and other mental events are routinely conceptualized as pushing symbols around (the Turing machine and its activities are essentially physical, and let us not forget that Newell and Simon's great conjecture about cognition is called the Physical Symbol System hypothesis). And we reason about nonphysical things such as wealth, information, and credibility in much the same way as we reason about physical things: We treat all of these things as resources like gasoline or food, to be produced, stored, consumed, traded, and so on. In sum, I think the dynamics of physical interactions with our environment is a solid foundation for concepts that represent physical and nonphysical activities, objects, relationships and attributes.

2. From Dynamics to Concepts

In this section I will develop a dynamical representation of verbs that denote physical interactions between two agents or objects named A and B. Examples include bump, hit, push, overtake, chase, follow, harass, hammer, shove, meet, touch, propel, kick, bounce, and so on.

I'll begin with some definitions. The *distance* between A and B, $D(A,B)$ is a projection of the not-necessarily physical locations of A and B onto a one-dimensional *progress space*. $P(A)$ and $P(B)$ are the locations of A and B in progress space and $D(AB) = P(B) - P(A)$. Note that the transformation of the states of A and B to $P(A)$ and $P(B)$ may be quite complex, and it might not even be physical. For instance, when a chef says he's "halfway done" with a meal, he is transforming the remaining tasks to a representation of the time required to finish the meal; this requires knowledge and skill. And when a professor asserts that a student is "advanced" relative to others she is mapping some attributes of the students to an entirely metaphorical line. For every domain, we must be able to map the "locations" of A and B (whether spatial coordinates or locations in a metaphorical space) into $P(A)$ and $P(B)$.

Velocities for A and B are defined in terms of $P(A)$ and $P(B)$, in the usual way, namely, $V(A) = dP(A)/dt$. Acceleration is just the derivative of velocity, $V'(A) = dV(A)/dt$. In physical space, relative velocity depends not only on $V(A)$ and $V(B)$, but also on the angle of A's trajectory relative to B's. In progress space, however, A and B are always traveling along a line. Since A and B are arbitrarily assigned labels, there are just four qualitative kinds of interactions between A and B in progress space:



In the first, A is behind B, and both are moving in the same direction; the point of contact is no closer than the rightmost agent and $D(AB) > 0$. In the second, A and B are moving toward each other in progress space and the point of contact is between them; again, $D(AB) > 0$. The third situation has A and B moving in the same direction, but their velocities are negative relative to the first situation, $D(AB) > 0$, and the point of contact is not closer than the leftmost agent. In the fourth situation, no contact can occur; I will not discuss this case any further.

In the first qualitative interaction, above, we define $V(A) \geq 0$ and $V(B) \geq 0$; in the second, $V(A) \geq 0$ and $V(B) \leq 0$. In the third, $V(A) \leq 0$ and $V(B) \leq 0$. We define relative velocity,

$$VR = V(A) - V(B).$$

For instance, if A's velocity is 10cm/sec. and B's is 20 cm/sec., but B and A are moving toward each other along a line (i.e., the second qualitative interaction, above), then $VR = V(A) - V(B) = 10 - (-20) = 30\text{cm/sec}$. In the third qualitative interaction, above, $VR = -30\text{cm/sec}$.

The interaction of A and B can be plotted in a two-dimensional space, called a *map*, as shown in Figure 2. (Maps are also called phase portraits, or phase diagrams; when the axes of a map represent values of a single variable measured at different times, the maps are called delayed coordinate embeddings. Some previous work in AI and Cognitive Science that uses maps as representations includes Rosenstein, et al, 1997; Bradley and Easley, 1997; Campbell and Bobick, 1995; Thelen and Smith, 1994) The horizontal dimension is $D(AB)$, the distance from A to B. The vertical dimension is VR , the relative velocities of A and B. The horizontal midline represents equal velocity, $V(A)=V(B)$. Above this midline, A is moving faster than B (or B is heading toward A, or both); below it, A is moving more slowly than B.

Some trajectories in this map are impossible. From the point labelled **a**, all trajectories must stay to the left of the vertical dashed line. This is because any vector from **a** to a point to the right of the line would mean A is slower than B but $D(AB) = P(B) - P(A)$ is decreasing. This can happen only if $P(A)$ is increasing faster than $P(B)$, which is

inconsistent with $V(A) < V(B)$. The shaded semicircle represents forbidden trajectories. Similarly, at point **b**, no vector can point left of the dotted line, because such a vector would represent B gaining on A (equivalently, A falling back toward B), which is inconsistent with A's velocity being higher than B's. At point **c**, the forbidden vectors flip from the left of the vertical line to the right, when A's velocity flips from being higher than B's to being lower.




Figure 2. Only some trajectories are physically possible





Figure 3. Some characteristic interactions between A and B

Point **d** illustrates that $D(AB)$ and velocities may change simultaneously. Imagine the vector to represent one time step of arbitrary duration. At the beginning of this interval, $P(A) = P(B)$ and B is moving faster than A. At the end of the interval, the velocities are equal but B is ahead of A.


The trajectory **e** shows five time steps of a "chase" behavior. In the first four steps, B is pulling away from A but at a decreasing rate, which is to say although A remains behind B, it speeds up relative to B, until, at the end of the fourth time step, the velocities are equal. At the end of the fifth time step, A's velocity exceeds B's, and A now starts to gain on B.

You can imagine that trajectory **e** is part of a clockwise, closed loop, as shown in Figure 3. Loops represent unending interactions in which B pulls away from A, then A gains on B, and so on. The loop entirely to the left of the $D(AB) = 0$ line in Figure 3 represents A's repeated failures to overtake B. The loop in Figure 3 that crosses the $D(AB) = 0$ line represents A and B "taking turns leading," like cyclists in a race. Finally, the open "spiral" that terminates at the point $D(AB) = 0, V(A) = V(B)$ begins with A and B at the same location, then has B pulling away rapidly, A catching up, and gently coming to rest at B.


This framework has sufficient representational power to describe many interactions between A and B, as shown in the following examples.



- a. $V(R)$ stays constant and relatively high until contact. "A runs into B full-tilt"
- b. VR decreases until contact: "touch, catch-up,"
- c. Looks like a "hit," as A speeds up as it approaches B
- d. "Drifting," barely moving toward each other because the relative velocity is nearly equal.



- a. Rapid deceleration, "hit the brakes."
- b. Initially A is losing ground to B, then "makes up for lost time," "comes storming back," "recoups its losses," "B eludes A briefly," etc.



- a. "B follows A, A leads B." Convoy, keeping close, etc.
- b. A and B are touching, either at rest or at matched velocities. Contact.
- c. "B narrowly escapes A" (because it started to move away from A very near the contact point)
- d. "B avoids A" (because a small effort, well before imminent contact, puts B out of reach for A).

Admittedly, some aspects of interactions between A and B are not represented. The directions of physical movement of A and B are not captured, only their relative positions in progress space (i.e., $P(A)$ and $P(B)$). Similarly, relative, not absolute velocities are represented. This means that the framework does not distinguish:

1. A and B are moving in the same direction and A is catching B because of superior velocity;
2. A and B are moving toward each other.


Hence, we cannot differentiate "A catches B" from "A and B embrace." Nor can we distinguish subtle intentional relationships between A and B. Suppose A and B are moving in the same direction, with B in the lead, and with $D(AB)$ varying in a narrow range. Is A trying to catch B while B tries to evade capture, or is A trying to follow B at a roughly constant distance? The dynamical maps I have described cannot represent this difference. However, I take up the subject of intentions in the next section.

An easily remedied representational deficit is that many verbs describe what happens when A and B make contact, whereas the previous examples all describe the interaction leading up to contact. Let us extend the framework to include types of contact.

2.1 Three Phases of Interactions


Physical interactions between agents can be viewed as having three phases. Consider the verb "push," for example. To push something, I first approach it and make contact with it. Generally, I try to achieve $VR = 0$ at $D(AB) = 0$, so that I gently touch the thing I'm trying to push. I apply force to it while remaining in contact for a period of time. When I or the thing I'm pushing breaks off contact, I may continue to move, or it may, or both. The three phases of a push, then, are **before**, **during**, and **after** contact. Many verbs of physical interaction can be represented in these terms; for example, a hit is like a push except that my velocity is high when I make contact and I stay in contact for a relatively short time. We have all received pushes that seemed a bit too much like hits; we might call them shoves. Where is the boundary between a push, a shove and a hit? There are no clear categorical boundaries: One's interpretation of an interaction depends on its dynamics, certainly, but also on contextual factors such as the intentions of the agents. I will return to intentions in the following section.

Once contact has been made, and a pair of agents is in the **during** phase, the salient dynamics concern position and energy exchange. Note that we don't care about relative position (i.e., distance between A and B) because by definition $D(AB)=0$ in the **during** phase. Similarly, relative velocity must be zero, otherwise relative position would change. A dynamic map for the during phase has the distance of the AB unit from the point of contact P_c on the horizontal axis, and the transfer of energy from A to B on the vertical dimension. We view the interaction from the perspective of agent A, and say $E(AB)>0$ if the net transfer of energy is to B, and $E(A,B)<0$ if B pushes harder.



- a. A transfers a lot of energy to B without any movement: A crashes into a brick wall (B).
- b. A transfers a lot of energy to B and the AB unit moves a little in the direction of A's movement. Pushing a car, a piano, or something else very massive.
- c. A initially transfers no energy to B, but ramps up to a constant flow, then ramps down. A pushes B.
- d. Like b except the AB unit moves in the direction of B's movement.

The denouement of the interaction between A and B is the **after** phase, which is entered when A and B break off contact. What seems most germane about this phase is the trajectories that A and B follow, so we could go back to the dimensions of **before** maps. A good reason to do so is that the **after** phase of one interaction may be the **before** phase of the next, especially for repetitive interactions such as tapping, hammering, harrassing, and so on:



- Both A and B remain at zero relative velocity and zero distance, attached.
- B's velocity with respect to A increases, as does its distance from A, then relative velocity goes to zero, and A and B remain at a constant distance. As if A kicked, shoved, shunted or otherwise provided impetus for B.
- Like b, except that A's velocity eventually increases again relative to B's, and the distance is reduced. This pattern would be observed in A hammering or harrassing B.
- A imparts some impetus to B, and B maintains it. "Kickstart, jumpstart, get B going, initiate B's action, etc."
- Like d except B keeps accelerating.
- Curiously, contact with B increases, rather than decreases A's velocity and thus its position relative to B. "slingshot, boost, accelerate," etc.
- Like f except achieving a constant relative velocity.
- A's velocity relative to B is apparently unaffected by contact. One imagines the **before** trajectory as the dotted line. This is what we'd expect to see if A overtakes B without making contact, or if B is insubstantial (e.g., fog) and offers no resistance to A.

Now let's look at some combinations of **before**, **during**, and **after** phases. Illustrative trajectories from each phase are shown in the three panels of Figure 4. Each trajectory in each panel has a label, and complete trajectories through the triptych are denoted by three-letter sequences. For instance, **cah** denotes A approaching B at a constant, high speed; contact for zero time with zero energy transferred (the black dot at the origin of the **during** phase); then A moving away from B at the same high speed. This trajectory represents "A overtakes B."




Figure 4. The **before**, **during** and **after** phases of physical interactions between A and B. The dashed vertical lines represent the point of contact, $D(AB)=0$. In the **before** and **after** phases, regions to the left of $D(AB)=0$ represent A behind B and regions to the right represent A ahead of B. In the **during** phase, regions to the right of $D(AB)=0$ represent displacement of the AB unit (remaining in contact) from the point of contact.

A remarkable number of verbs can be represented in this framework:

- aaa** A approaches B, touches it, and remains in contact with it. *A gently touches B* with no net transfer of energy between them. Relative velocity is inherently ambiguous: We know A and B have equal velocities in the **after** phase, but we don't know whether this velocity is zero.
- ada** A approaches B, makes contact, then gradually increases the energy it transfers to B, maintains a level of energy transfer, then ramps down. A and B remain in contact in the **after** phase. *A pushes B*.
- adb** A approaches B, makes contact, and gradually (d) or rapidly (c) increases the energy it transfers to B. In the **after** phase, B moves a little ahead of A.
- acb** Initially its velocity increases relative to A's then decreases. Depending on the rate of energy transfer, the amount transferred, and the distance B moves in the **after** phase, this is *kick, nudge, shove, propel*, and so on.
- The movement in the **before** phase is inherently ambiguous: We don't know whether A is moving toward B, B is moving toward A, or both. Similarly, the increasing distance between A and B in the **after** phase might occur because A stops moving (or slows down) but B continues, or because B stops and A is recoiled, or a combination of effects. Thus, **acb** represents *A bounces off B* as well as *kick, shove*, and so on. Similarly, **acb** represents *symmetric repulsion*, where A and B approach each other, make contact, then bounce away from each other.
- aca** As above, except B doesn't move. Depending on rates and amounts of energy transferred, this too may be a *kick* or a *bump* (but not a *shove* or *propel*, because B doesn't move). Alternatively, **ada** denotes a more gradual interaction, as in *A leans against B, A strains against B*.
- bce** Whereas **b** in the **after** phase represents A and B moving apart with an increasing, then decreasing, velocity, trajectory **e** represents A and B moving apart with a strictly *increasing* velocity. Imagine a hand (A) pushing a cup (B), off the edge of a table. Or we might say *A dislodges B*, or *frees* it from some stricture. Or B might *flee* from contact with A.
- cba** A and B converge at a high, constant rate. At the instant of contact they exchange a lot of energy, and remain in contact during the **after** phase. This is what happens when a car *crashes* into a tree. More benignly, B may absorb all A's energy with no ill effect, but I know no verb to describe this interaction.
- dcc** This is a cyclic interaction where A and B converge, energy is transferred, and during the **after phase**, A and B diverge then converge again. Many verbs denote this repetitive pattern: *Hammer, harass, clap*, and so on.
- bbf** A accelerates relative to B until the point of contact, B absorbs energy from A, and A is slowed down and eventually comes to rest a little beyond B. *A pushes through B*.
- bbg** Like **bbf**, except A maintains a constant velocity after interacting with B. *A breaks free of B*.

As with the individual maps, this triptych represents many aspects of interactions but fails to represent others. Some ambiguities have already been discussed (see **adb** and **acb**, above). Because this framework doesn't represent actual spatial coordinates, it cannot differentiate the cases in Figure 5. Similarly, we cannot tell whether A is pushing an unyielding B, or A and B are pushing against each other. Another source of ambiguity arises because the framework doesn't specify what kinds of things A and B are. In particular, it is unclear what kind of energy A transfers to B and where this energy comes from. If A transfers kinetic energy, then the sequence **ab...** will in some cases be physically impossible because once the relative velocity of A and B reaches zero, there is no kinetic energy to transfer. On the other hand, if A is capable of generating movement itself, as most agents are, then it can transfer kinetic energy to B even after their velocities are matched, simply by increasing its velocity. Another ambiguity arises because no scales are specified in the maps. We can say one interaction involves more force than another (e.g., a shove versus a tap), but if we have only one trajectory and cannot calibrate it against others, then we cannot judge whether it is gentle or violent.

Despite these and other limitations in representational power, the framework is extremely compact and simple, yet captures a very large number of verb meanings. Finer distinctions in meaning can be had by adding dimensions to the maps (e.g., x and y spatial dimensions). There is obviously a tradeoff between the expressivity and complexity of the maps, but I find it remarkable that a representation this simple is so expressive.




Figure 5. The framework cannot differentiate cases in which A and B recoil mutually; or B is propelled but A doesn't move; or A bounces off B; or A and B bounce off each other at unspecified angles.

3. Intentions

Suppose at a party a drunk man pats you on the back a little too hard, knocking you forward. Is this a pat gone awry, or a not-too-subtle aggression? You don't know. You don't know his intention. Figure 6 shows two representations of the interaction. The actual trajectory is the same in both: His hand makes contact with your back at a relative velocity greater than zero, and it transfers a considerable amount of energy to your back. The difference between the representations is the man's *goal regions*. On the left, you see a benign pat gone wrong. The goal region for relative velocity (the shaded area in the **before** phase) is considerably lower than the one the man actually generated (he's drunk, after all). The trajectory for energy transfer and displacement of your back falls well outside his goal region, also. On the right, however, the man generates the relative velocity profile that he intends, and he hits you as hard as he intends, and you are

knocked forward as far as he intends (pretty accurate for a drunk!) This is a malicious act of violence.




Figure 6. Goal regions encode intentions

Generally, you don't know the intentions of other parties, so you cannot be sure that "hit" is the correct verb to describe an interaction like this. But sometimes, dynamics alone are sufficient to infer intent, albeit heuristically. If the man doesn't intend to hurt you, then he will probably try to modulate his arm movement when he realizes it is too fast. If he appears to try to check his swing, fine, but what if he actually increases his arm speed as he approaches you? Then it requires uncommon charity (or stupidity) to excuse his behavior as accidental.

Another example involves repetitive behavior. Suppose you are on the highway and indicate to change lanes. As you start your move, the car in front does, too, blocking you. Fine. It could be accidental. So you try again, and again the car in front blocks you. After a couple of interactions like this, you strongly suspect that the driver in front has a goal region for your mutual interaction, and the goal region has him in front. In Figure 4, his goal region would be shaded to the left of the point of contact: You should always be behind him. In this case, it would be appropriate to use verbs like *block*, *obstruct*, *impede*, and so on to describe the interaction. These are intentional verbs, but as I have indicated, intent can sometimes be inferred from dynamics.

Naturally, if you can label a map with the intentions of both parties, then you can use even more specific verbs. If B wants to keep A behind him, and A wants to be in front, then we can say A is trying to *escape*, *slip by*, *get away from*, and so on.

4. Learning and Reasoning

The framework described here has been implemented in part, and we have run experiments to see how accurately a system can reason with these representations and how easily it can learn them. Some of these experiments are reported elsewhere (Rosenstein et al., 1997) and some are ongoing. I will summarize the results here.

Michael Rosenstein and I studied the **before** phase of the interactions between two simulated agents (Rosenstein et al., 1997). Each agent exhibited one of nine behaviors. For instance, A might try to *avoid* B while B tries to *hit* A. We discovered that a system can very quickly learn the maps associated with each of the 81 pairs of behaviors. It could then use a trajectory – a time series of x,y locations of A and B – to recognize the

appropriate map with high accuracy, and once it had identified the map it could predict accurately the outcome of the interaction (avoidance, contact, or a perpetual chase). The learning in this case was supervised in the sense that we told the system which of 81 pairs of behaviors it was observing, and it merely learned the dynamics of the interaction. Recently we have developed an unsupervised version, where the system clusters training trajectories together without knowing which behaviors generated them. Remarkably, with very little training, the system came up with six clusters: Three represent types of interaction where A escapes B. The first kind of escape is the case where B never gets close to A. The second is the case where B nearly reaches A, but A slips away. The third is the case where B's momentum causes it to overshoot A, which escapes. The fourth cluster represents cases where B catches A. The fifth and sixth clusters represent versions of perpetual chasing.

What is remarkable about these results is that we did not tell the system to cluster trajectories by their outcomes, but time series are so redundant that clustering by trajectory dynamics produces clusters that have qualitatively different outcomes. In short, simply clustering trajectories seems sufficient to produce qualitatively different classes of interactions. We are well on our way to learning classes and concepts through physical interaction, without supervision.

The only reasoning our system currently does with maps is recognition of trajectories and prediction of outcomes of interactions. The heuristic reasoning about intentions, described in the previous section, has not been implemented. Even so, recognition and prediction are powerful modes of reasoning when interacting with moving objects.

5. Conclusion

I have introduced a representation for the meaning of verbs based on dynamics. It is a simple but remarkably expressive representation, and its expressivity can be augmented very easily by adding dimensions and goal regions. Experiments suggest that a system can quickly learn to recognize interactions based on this representation, and to predict their outcomes. Intriguingly, unsupervised learning produces clusters of interactions that have qualitatively different dynamics and outcomes, suggesting that dynamics and clustering may be all one needs to learn classes of interactions.

6. References

- Bradley, L and Easley, M. 1997. *Reasoning about sensor data for automated system identification*. In *Advances in Intelligent Data Analysis*. X. Liu, P. R. Cohen and M. Berthold, Eds. Lecture Notes in Computer Science, 1280. Springer. pp. 561–572.
- Campbell, L. and Bobick, A. 1995. *Recognition of human body motion using phase space constraints*. MIT Media Laboratory Perceptual Computing Section Technical Report No. 309.

- Elman, J. 1995. *Language as a dynamical system*. In *Mind As Motion*, R. F. Port and T. van Gelder, Eds. MIT Press. pp. 196 – 225.
- Lakoff, G. 1984. *Women, Fire and Dangerous Things*. University of Chicago Press.
- Lakoff, G. and Johnson, M. 1980. *Metaphors We Live By*. University of Chicago Press.
- Regier, T. 1995. A model of the human capacity for categorizing spatial relationships. *Cognitive Linguistics*, 6 (1), pp. 63 – 88.
- Rosenstein, M, Cohen, P. R., Schmill, M. D., and Atkin, M. S. 1997. *Action, representation, prediction and concepts*. Presented at the 1997 AAAI Workshop on Robots, Softbots, Immobots: Theories of Action, Planning and Control. Also available from <http://eksl-www.cs.umass.edu:80/publications-2.html>
- Thelen, E. and Smith, L.B. 1995. *A Dynamic Systems Approach to the Development of Cognition and Action*. MIT Press.